

A

- a option
 - for cp, 69
 - for ifconfig, 56, 124
 - for uname, 168
- Access control lists (ACLs), 95
- Address Resolution Protocol (ARP)
 - for Ethernet, 11, 58
 - for host setup, 89
- Address space manipulation, 199
- Addresses. *See* IP addresses; MAC addresses
- Administration in captive UML, 287–289
- Administrators, console access by, 224–225
- aio_abi.h file, 249
- AIO facility, 192–193
- aio option, 304
- allmodconfig configurator, 239
- allnoconfig configurator, 239
- Always disallow TCP connections to X server
 - option, 96
- anon driver, 227–228
- Apache servers, 94
- append switch for hostfs, 215–216
- Application administration in captive UML, 287–289
- Application-level clustering, 289–294
- Application programming interface in captive UML, 289
- ARCH, 239
- ARP (Address Resolution Protocol)
 - for Ethernet, 11, 58
 - for host setup, 89
- arp command
 - for network interfaces, 255–257
 - for TUN/TAP, 127–128, 131
- Attacks
 - humfs for, 216
 - packet faking, 122–123
 - with TUN/TAP, 130
- Audio pass-through driver, 301

Authentication

- in captive UML, 288
- in MConsole requests, 185

Authorization

- in captive UML, 288
- in host setup, 95
- in MConsole requests, 185

B

- b command in sysrq, 173
- b switch
 - for tunctl, 124, 311
 - for uml_moo, 71, 308
- Backing files
 - COW, 62–65
 - merging with, 70–71
 - moving, 69–70
- Backups
 - COW files for, 64
 - for filesystems, 116–117
- bash command, 57
- Bind mounts, 214–215
- Block devices
 - configuring, 170, 303
 - pluggable, 87
 - using and abusing, 83–87
- Block drivers, 23–25
- Booting
 - clusters, 268–272
 - from COW files, 67–68
 - first time, 20–24
 - successful, 24–28
- Bottlenecks, 203, 208
- brectl utility, 137–138
- Breakpoints, 178
- bridge-utilities package, 137
- Bridging
 - security in, 140
 - setting up, 136–139

- Broadcast domains for host setup, 89
- BSD jail, 2
- Buffered I/O, 194
- Bug fixes, 234
- Bug-tracking system, 287
- Builds, 249–250
- BusyBox project, 10
- bzip files, 87

C

- c switch for ubd, 303
- Cached data, 115
- cad command, 310
- CAP_SYS_RAW, 213
- Capabilities, permissions for, 213
- Captive UML, 283
 - application administration, 287–289
 - application-level clustering, 289–294
 - Evolution, 286–287
 - secure mod_perl, 283–285
 - standard application programming interface, 289
- Carvalho de Melo, Arnaldo, 7
- chroot technology, 2, 71, 216–217, 220–221
- CKRM (class-based linux Kernel Resource Management), 297
- Clock
 - real-time, 245–246
 - synchronizing, 28
- close calls, 113
- cluster.conf file, 267–269
- Clusters, 265–268
 - application-level, 289–294
 - available, 273
 - booting, 268–272
 - exercises, 272–273
- CMDLINE_ON_HOST option, 241–242
- Code pages, 201
- Collaboration, clustering for, 292–293
- Command-line options, 301
 - debugging, 303–304
 - device and hardware specifications, 301–303
 - informational, 305
 - management, 304–305

- Commands, running within instances, 180–182
- Compartmentalization, 297–298
- Compiling, 233–234
 - builds, 249–250
 - configuration interfaces, 235–240
 - configuration options
 - console, 248
 - debugging, 249
 - execution mode-specific, 240–243
 - generic, 243–246
 - networking, 247–248
 - virtual hardware, 246–247
 - source downloading for, 234–235
- CON_CHAN option, 248
- CON_ZERO_CHAN option, 248
- config command
 - for devices, 169
 - process context for, 186
 - in uml_mconsole, 310
- config configurator, 237–238
- CONFIG_EXTERNFS option, 213–214
- CONFIG_HOSTFS option, 213–214
- CONFIG_IP_MROUTE option, 152
- CONFIG_IP_MULTICAST option, 152
- CONFIG_MAGIC_SYSRQ option, 199
- CONFIG_MODE_SKAS option, 199
- CONFIG_MODE_TT option, 199–200, 218
- CONFIG_STATIC_LINK option, 218
- configfs filesystem, 266
- "connection refused" message, 96
- Connectivity with TUN/TAP devices, 125–129
- Consistency problem, 115
- Consoles, 40–47
 - configuring, 170, 248
 - for host ports, 41–45
 - MConsole. *See* SKAM Management Console (MConsole)
 - security for, 223–225
- Consolidating servers, 8–10
- Contexts
 - forcing threads into, 177–179
 - process, 186, 231
- Cookies, Xauthority, 95
- Copy-On-Write files. *See* COW (Copy-On-Write) files

- Copying data into instances, 83–87
- Corrupted filesystems, 68
- Cost savings, 8
- COW (Copy-On-Write) files, 11, 61–66
 - backing files for, 62–65
 - merging with, 70–71
 - moving, 69–70
 - for backups, 117
 - booting from, 67–68
 - in small server setup, 208
 - sparseness of, 175
- cp command, 69
- cpuinfo file, 36, 219
- CPUs, multiple, 243–244
- Ctrl-Alt-Del handler, 171–172
- Ctrl-C, signals from, 203
- Ctrl-Z, signals from, 203

D

- d switch
 - for screen, 205
 - for tunctl, 124, 311
 - for ubd, 303
 - for uml_moo, 71
- Daemon transport, 152, 154
- Databases
 - in captive UML, 288
 - metadata, 113–115
 - package, 278–282
- date command, 70
- dd command
 - for copying data into instances, 85–87
 - for copying files, 53
 - for swap space, 47
- debug option, 304
- Debugging, 13
 - options for, 249, 303–304
 - PT_PROXY for, 242
- Decision-making for hardware, 273–274
- Default Apache install page, 94
- Default configuration, 239–240
- Default gateways for uml_switch, 163
- Default ports for multicast transport, 152
- Default routes
 - for host setup, 91
 - for TUN/TAP, 133

- Default values for transports, 148
- defconfig configurator, 235, 239
- Deleting routes, 77
- Denial-of-service attacks, 216
- dev with jails, 219
- Development uses, 12–13
- Devices
 - hardware specifications for, 301–303
 - memory-mapped I/O for, 82
 - queries for, 169–170
 - TUN/TAP. *See* TUN/TAP devices
- devpts filesystem, 34
- df command, 34
- DHCP
 - in bridging, 139–140
 - for host setup, 89
 - for transports, 148
 - through TUN/TAP devices, 134–135
- dhcp-fwd service, 134–135
- dir_metadata file, 112
- Disable option for xconfig, 236
- Disaster recovery, 13–14
- Disk numbers for partitions, 50
- Disks
 - listing, 35–36
 - partitioned, 49–52
 - as raw data, 53–54
 - saving space on
 - COW files for, 62–66
 - humfs for, 111
 - for swap space. *See* Swap space
- diskstats file, 261
- DISPLAY environment variable, 96–97
- dmesg command, 22
 - for consoles, 40, 42
 - for host setup, 90
- DocumentRoot, 94
- Downloading source, 234–235
- Drivers
 - block, 23–25
 - initializing, 22–23
 - loopback, 24–25
- DSL connections, 136
- dsp option, 301
- Duality of UML, 18
- Dumping
 - memory statistics, 172–173
 - registers and stack, 174, 178

Dynamic linking
 configuration option for, 240–241
 with jails, 218

E

e command, 173
 -e switch for xterm, 302
 e2fsprogs-devel package, 266
 EAGAIN value, 36
 ebttables, 140–141
 Educational uses, 10–12
 Efficiency, filesystem, 119
 Emacs, 289
 Embedded hardware, 13
 Emulating devices, 82
 Enable option for xconfig, 236
 Encapsulation in SLIP, 144
 ERR message, 182
 Error indicators in MConsole requests, 185
 eth option, 302–303
 Ethernet
 in bridging, 136, 139
 in host setup, 89
 for instances, 72–73
 IP addresses for, 57–58
 in SLIP transport, 144
 Ethernet cards, 134
 Ethertap
 configuring, 150
 for frames, 54
 for host network access, 143
 for host setup, 88
 Evolution, 286–287
 Exchanging packets, 72–73
 Exclusive locks, 68
 exec command, 181
 Execution modes
 options for, 240–243
 in small server setup, 194–196
 patches for, 201–202
 skas0, 200–201
 skas3, 198–200
 tt, 197–198
 Vanderpool and Pacifica, 202–203
 ext2 filesystem, 24, 118

ext2online filesystem, 118
 ext3 filesystem, 118
 Extending filesystems, 117–118
 externfs filesystem, 277–282

F

f command in sysrq, 173
 -f switch
 for log, 175, 310
 for tunctl, 311
 fake_ide option, 303
 Faking packets, 122–123
 fd directory, 31
 fd file descriptor, 43–44
 fdisk tool, 50–51
 file command, 19, 24
 file_metadata file, 109, 112
 Filenames for backing files, 69
 Filesystems, 101
 backups for, 116–117
 booting, 26
 corrupted, 68
 extending, 117–118
 externfs, 277–282
 host access to, 114–116
 host directory mounting, 101–104
 with hostfs, 104–108
 with humfs, 108–114
 selecting, 119–120
 filesystems file, 34–35, 102–103
 Filters
 for pcap, 154
 for TUN/TAP, 130
 Firewalls
 in host setup, 92–93
 for TUN/TAP, 132
 fonts-xorg-75dpi package, 96n
 Forcing threads into contexts, 177–179
 FORWARD chains, 141
 Frames
 in bridging, 136, 139
 host setup for, 88–89
 transmission of, 54
 free command, 52
 fsck message, 26

fstab file

- for small server setup, 207
- for swap space, 52
- sync options in, 115

Future of UML, 14–15, 275–277

- captive UML, 283
 - application administration, 287–289
 - application-level clustering, 289–294
 - Evolution, 286–287
 - secure mod_perl, 283–285
 - standard application programming interface, 289
- conclusion, 298–299
- externfs filesystem, 277–282
- virtual processes, 282–283
- virtualized subsystems, 295–298

fwwm window manager, 97**G****gconfig configurator, 238****gcov, 12, 249****GCOV option, 249****gdb**

- in debugging, 12, 178
- with ptrace, 242

gdb-pid option, 304**gdmsetup, 96****getpid loops, 226****gettimeofday command, 28, 245****getty**

- for consoles, 45–46
- for virtual serial lines, 79–82

gettys, 29**GFS clusters, 273****Giarrusso, Paolo, 7, 200, 231****GID (group ID) root, 109****GIMP, 289****glibc, 192****glibc2-devel package, 266****go command, 310****gprof, 12, 249****GPROF option, 249****Group ID (GID) root, 109****GRUB command, 23****GTK toolkit, 238****Guest scheduler, 295–297****H****halt command**

- vs. cad, 171
- for instances, 169
- process context for, 186
- for shutdown, 59
- in uml_mconsole, 309

Hang notifications, 186, 188**Hardware**

- configuration options for, 246–247
- decision-making for, 273–274
- developing, 13
- queries for, 169–170
- specifications for, 301–303

Header files, 249**help command, 179–180, 309****--help option, 305****HighFree field, 258****HIGHMEM option, 244****Highmem support, 258**

- in skas mode, 200

- in small server setup, 208

- in tt mode, 197–198

HighTotal field, 258**History of UML, 4–8****home, mounting, 214****HOST_2G_2G option, 241****Host directory mounting, 101–104**

- with hostfs, 104–108
- with humfs, 108–114

host filesystem, 103**hostfs, 101–102**

- advantages of, 119
- append switch, 215–216
- for bind mounts, 215
- for file access, 114–116
- for host directory mounting, 104–108
- for mount restrictions, 214

Hosts

- consoles for, 41–45
- filesystem access by, 114–116
- instance management from. *See* Instances; Management Console (MConsole)
- intercepting and nullifying calls to, 226
- kernel in, 18
- memory consumption by, 25–26
- in networking, 87–99, 143–145

Hosts *continued*
 proxy arp for, 58
 for serial lines, 79–81
 Hot-plug memory, 228–230
 Hot-plugging devices, 169–170
 httpd file, 94
 httpd.conf file, 94, 285
 https sessions, random numbers for, 246
 -hub switch for `uml_switch`, 153, 311
 Hubs, 72
 humfs, 101–102, 104
 advantages of, 119
 for denial-of-service attacks, 216
 for file access, 114–116
 for host directory mounting, 108–114
 in small server setup, 208
 humfsify command
 for humfs, 110–111
 reference, 307–308
 hwclock program, [28](#)
 hwrng file, 247

I

`i` command in `sysrq`, 173
`ifconfig` command, 54–56
 for host setup, 90
 for instances, 73–77
 for interfaces, 124, 253–254
 Informational options, 305
 Inheritance of capabilities, 213
`init` process, 171
 Initializing drivers, 22–23
`initrd` option, 302
`inittab` file
 for Ctrl-Alt-Del handler, 171
 editing, 45–46
 for serial lines, 40, 79–80, 82
 INPUT chains, 141
 Instance kernel log, 175
 Instances
 block devices for, 83–87
 for console server, 224–225
 COW files for. *See* COW (Copy-On-Write) files
 halting and rebooting, 169, 171

jailing, 216–223
 managing, 167
 with Management Console. *See* Management Console (MConsole)
 with signals, 188–189
 networking, 71–79
 running commands within, 180–182
 sending interrupts to, 179
 stopping and restarting, 174–175
`int` command, 309
 Intercepting host system calls, 226
 Internal utilities, 312
 Internet collaboration, clustering for, 292
 Interprocess communication (IPC)
 mechanisms, 72
 Interrupts
 handling, 186
 sending to instances, 179
 interrupts file, [35](#)–[36](#)
 I/O
 AIO facility for, 192–194
`MADV_TRUNCATE` for, 228
`iomem` driver, 247, 302
 IP addresses
 in bridging, 137, 139
 for Ethernet, 57–58
 for hosts, 87–90, 93, 96
 for instances, 73, 75–76
 reusing, 56–57
 for transports, 148
 for TUN/TAP, 124–126, 130–134
 for virtual serial lines, 82
 IPC (interprocess communication)
 mechanisms, 72
 iptables
 for bridging, 142
 for filtering, 130–132
 for host setup, 92–93
 Isolated networks, transports for,
 145–146

J

jail switch
 for `hostfs`, 215
 for `uml_moo`, 71

Jailed processes, 2–3
Jailing instances, 214–224

K

Kernel
 logs for, 173
 logging to, 175
 for multicast networks, 155
 security for, 212–214
 versions of
 queries for, 168–169
 in small server setup, 192–194
 virtualized subsystems in, 295
KERNEL_HALF_GIGS option, 242–243
Kernel-level programming, 12
Kernel mode, 212
Kernel modules, 212–213
KERNEL_STACK_ORDER option, 244
Kernel tree, 234–235
Keyboards listing, 35–36
Killing tasks, 173
Kroah-Hartman, Greg, 234

L

Large numbers of devices, configuring
 memory, 257–265
 network interfaces, 252–257
Large server management, 211
 final points, 232
 future enhancements
 MADV_TRUNCATE, 227–230
 PTRACE_FAULTINFO, 227
 remap_file_pages, 230–231
 sysemu patch, 226–227
 VCPU, 231
 security for
 configuration for, 212–216
 console, 223–225
 jailing instances, 216–223
 skas3 vs. skas0, 225–226
LDT (Local Descriptor Table) entries, 199
len field in MConsole requests, 185
lib file, 217
libpcap, 145

Libraries with jails, 217–218
LILO command, 23
Linking, dynamic and static
 configuration option for, 240–241
 with jails, 218
Links, symbolic, 113
linux file, 249–250
Local Descriptor Table (LDT) entries, 199
Locks
 in application-level clustering, 291
 for instances, 68
log command, 175, 310
log level setting, 173
Login prompt, 29
Logins
 console for, 224
 as normal users, 39–40
 for running commands, 181
Long-lived instances, 203–205
longjmp command, 179
Loop-mounting images on hosts, 115
Loopback drivers, 24–25
LowFree field, 258
LowTotal field, 258
ls command line, 47
ltrace, 12
Lustre clusters, 273

M

m command in sysrq, 172–173
-m switch for screen, 205
MAC addresses
 in bridging, 136, 139
 for host setup, 90
 for instances, 75
 for transports, 148
 for TUN/TAP, 127, 134–135
MADV_TRUNCATE patch, 227–230
Magic SysRq facility, 116–117
MAGIC_SYSRQ option, 246
Management Console (MConsole)
 for backups, 117
 MConsole protocol, 183–186
 notifications, 186–188
 for partitions, 50
 Perl library, 185

- Management Console (MConsole) *continued*
 - for queries. *See* Queries, MConsole
 - requests in, 184–186
 - uml_mconsole client, 182–183
 - Management options, 304–305
 - Mapping
 - file operations to host operations, 104–108
 - memory, 230–231, 244
 - in skas3, 199
 - Masquerading, 92–93
 - Master UMLs, 80–81
 - Maximal Transfer Units (MTUs), 136
 - mcast command, 89, 151–152
 - MConsole. *See* Management Console (MConsole); Queries, MConsole
 - mconsole driver, 23–24, 304
 - MCONSOLE option, 246
 - MConsole protocol, 183–185
 - mem file
 - for kernel access, 213
 - for swap space, 48
 - mem option, 302
 - meminfo command, 261–262
 - for debugging, 176
 - for instances, 258
 - output from, [30–32](#)
 - for scalability limits, 252
 - for swap space, 48
 - Memory
 - configuring, 170, 257–265
 - consumption of
 - host, 25–26
 - monitoring. *See* meminfo command
 - Highmem support for, 197–198, 200, 244
 - mapping, 230–231, 244
 - saving
 - COW files for, 62, 64, 68
 - MADV_TRUNCATE for, 227–230
 - small server setup for, 206–208
 - statistics dumping for, 172–173
 - swap space for. *See* Swap space
 - usage information, [30–32](#)
 - Memory-mapped I/O, 82
 - Memory pages, 198, 230–231
 - menuconfig configurator, 236–237
 - Merging COW files with backing files, 70–71
 - metadata file, 112
 - Metadata for files, 109–115
 - mixer option, 301
 - mkfs for clusters, 270
 - mm process, 199, 201–202
 - mmap, 104, 201, 208
 - MMAPPER option, 247
 - mnt directory, 105
 - mod_perl module, 283–285
 - mode option, 304
 - MODE_SKAS option, 240
 - MODE_TT option, 240
 - Modification time for backing files, 69–70
 - modprobe command, 57
 - Modular option for xconfig, 236
 - Molnar, Ingo, 230–231
 - Monitoring memory consumption. *See* meminfo command
 - Morton, Andrew, 7, 231, 265
 - mount command, 105
 - Mounting host directories, 101–104
 - with hostfs, 104–108
 - with humfs, 108–114
 - Moving backing files, 69–70
 - mprotect, 201
 - MTUs (Maximal Transfer Units), 136
 - Multicasts, 72–73
 - configuring, 151–152
 - example networks, 155–160
 - with instances, 75–78
 - for isolated networks, 145
 - Multiple clustered applications, 292
 - Multiple instances, COW files for. *See* COW (Copy-On-Write) files
 - Multiple processors, 243–244
 - Multiple users, hostfs with, 107
 - munmap, 201
- ## N
- n command in sysrq, 174
 - n option for uname, 168
 - Name server responses, faking, 123
 - Named pipes
 - hostfs with, 108
 - humfs with, 110
 - Names
 - for devices, 124

- for partitions, 50
 - for screen sessions, 205
- nepus option, 302
- NEST_LEVEL option, 241
- Network Address Translation (NAT), 92
- network file, 269
- Network sniffers, 146
- Networking, 54–59, 121
 - configuration options for, 247–248
 - examples
 - multicast, 155–160
 - summary, 166
 - uml_switch, 160–166
 - filesystem access in, 104
 - hosts in, 87–99
 - instances, 71–79
 - interface configuration for, 170, 252–257
 - manual setup for
 - bridging, 136–142
 - TUN/TAP. *See* TUN/TAP devices
 - small server setup for, 206
 - transports, 142–143
 - configuring, 147–154
 - for host network access, 143–145
 - for isolated networks, 145–146
 - selecting, 146–147
- New connection message, 162
- NFS clusters, 273
- nfs directory, 103
- Nodes, cluster, 268–272
- nodev entries, 102
- nolisten tcp, 96
- Nonbroadcast frames in bridging, 139
- Nondevice filesystems, 102
- none device, 43
- Nonexclusive read-only locks, 68
- nooptimize flag for pcap, 154
- noprocm option, 304
- noprodefaultinfo option, 304
- Normal user logins, 39–40
- nosyemu option, 304
- Notifications
 - for jails, 222–223
 - MConsole, 186–188
- NR_CPUS option, 243–244
- NULL_CHAN option, 248
- null device, 43
- Nullifying host system calls, 226

O

- O_APPEND option, 215
- O_DIRECT I/O
 - caches in, 119
 - in host kernel, 192–194
 - in small server setup, 208
- o option for hostfs, 105–106
- o2cb file, 267
- ocfs2 clusters, 290–291, 293–294
- ocfs2 script, 265–267
- ocfs2console, 266
- od utility, 86
- OK message, 182
- oldconfig configurator, 235, 238
- Omitted transport parameters, 148
- open calls, 113
- Openswan project, 10
- optimize flag for pcap, 154
- Oracle, 265
- Out-of-memory condition, 173
- OUTPUT chains, 141
- Outside network access, 132–133
- Overwriting files, preventing, 215
- Ownership of files, 106–110, 112

P

- p command in sysrq, 174
- p switch for cp and tar, 69
- Pacifica execution mode, 202–203
- Packages, databases for, 278–282
- Packets
 - exchanging, 72–73
 - faking, 122–123
 - forwarding, 127
 - with instances, 74, 77–78
 - transmission of, 55
- Page-by-page memory mapping, 230–231
- Page faults, 199–200
- Panic notifications, 186, 188
- Parameters for transports, 148
- Partitioned disks, 49–52
- passwd file, 101
- password prompt for running commands, 181
- Passwords in captive UML, 288

- Patches, 234
 - for execution modes, 201–202
 - for performance, 226–227
- pcap transport
 - configuring, 154
 - for isolated networks, 145–146
- Performance
 - bottlenecks in, 203, 208
 - COW files for, 64
 - memory for, 259–265
 - PTTRACE_FAULTINFO patch for, 227
 - remap_file_pages for, 230–231
 - in skas3 Mode, 198
 - in SLIP transport, 144
 - in small server setup, 208
 - sysemu patch for, 226–227
- Perl library, 185
- Permissions
 - in application-level clustering, 291
 - for capabilities, 213
 - for files, 106–110, 112
 - for host setup, 95
 - for security, 123
- physdev module, 142
- Physical memory, small server setup for, 206–208
- pid file with jails, 222–223
- PIDs (process IDs) for signals, 189
- ping command, 58
 - for bridging, 141
 - for host setup, 90–91, 93–94
 - for instances, 76–78
 - for multicast networks, 158–160
 - for network interfaces, 254, 256–257
 - for TUN/TAP, 125–129, 133
 - for uml_switch, 162–166
- Pipes
 - with hostfs, 108
 - with humfs, 110
 - with uml_switch, 153
- Pluggable block devices, 87
- Point-to-Point Protocol (PPP)
 - for frames, 54
 - for host setup, 88
- PORT_CHAN option, 248
- port device, 43
- port-helper utility, 312
- Ports
 - consoles for, 41–45
 - for multicast transport, 152
 - with Slirp, 144
 - for uml_switch, 162
- PPP (Point-to-Point Protocol)
 - for frames, 54
 - for host setup, 88
- ppp0 device, 92
- PPPoE connections, 136
- print statement, 13
- Privileged contexts, 231
- Privileges
 - with jails, 220
 - in virtual machines, 9
- proc command, 176
 - for files, 310
 - process context for, 186
- proc directory
 - for cpu, ~~36–37~~, 219
 - for diskstats, 261–264
 - examining, 176
 - for filesystems, ~~34–35~~, 102–103
 - for guest scheduler, 295
 - for interrupts, ~~35–36~~
 - for mconsole, 187–188
 - for memory. *See* meminfo command
 - for mm, 199, 201
 - for sysrq, 172
- Process contexts, requests in, 186
- Process IDs (PIDs) for signals, 189
- process_kern.c file, 178
- Processes, 18
 - contexts for, 231
 - in execution modes. *See* Execution modes
 - jailed, 2–3
 - listing, 29–30
 - permissions for, 213
 - virtual, 282–283
- Processors, multiple, 243–244
- procfv filesystem, 34
- promisc flag for pcap, 154
- Proxies for performance, 259
- Proxy arp
 - for host routing, 58
 - for TUN/TAP, 128
- ps command, ~~29–30~~

Pseudo-terminals, 79–81
 PT_PROXY option, 242
 ptrace
 gdb with, 242
 for intercepting system calls, 231
 in skas3 mode, 202
 in sysemu, 226
 in tt mode, 199
 PTRACE_FAULTINFO patch, 200–202, 227
 PTRACE_LDT option, 202
 PTRACE_SWITCH_MM option, 199
 pts device
 for consoles and serial lines, 42–44, 79–81
 with jails, 218–219
 PTY_CHAN option, 248
 pty device, 43
 Pulavarty, Badari, 227

Q

Queries, MConsole, 168
 for Ctrl-Alt-Del handler, 171–172
 for forcing threads into contexts, 177–179
 for halting and rebooting instances, 169, 171
 for hardware configuration, 169–170
 for help, 179–180
 for logging to instance kernel log, 175
 for proc, 176
 for running commands, 180–182
 for sending interrupts, 179
 for stopping and restarting instances, 174–175
 for SysRq handler, 172–174
 for version, 168–169
 quit command, 309
 Quotas on hosts, 118

R

-r switch
 for screen, 204–205
 for ubd, 303
 for uname, 168
 randconfig configurator, 238
 random file, 247

Random numbers, 246–247
 Raw data, disks As, 53–54
 Read-only files, 64
 Read-only locks, 68
 Read-write locks, 68
 Reading files, 113
 readlinedevel package, 266
 Readlinks, 113
 Real-time clock, 245–246
 Real-time tasks, 174
 reboot command
 vs. cad, 171
 for instances, 169
 process context for, 186
 in uml_mconsole, 310
 Rebooting instances, 169, 171
 Registers, dumping, 174
 remap_file_pages call, 230–231
 Remote logins, 94
 remove command
 for devices, 169
 process context for, 186
 in uml_mconsole, 310
 Requests
 MConsole, 184–186
 web site, 284–285
 Resizing filesystems, 118
 resolv.conf file
 for host setup, 91–92
 for TUN/TAP, 132–133
 respawn command, 82
 Restarting instances, 174–175
 Restoring timestamps, 70
 root option, 302
 Root privileges
 and capabilities, 213
 with jails, 220
 in virtual machines, 9
 /rootfs switch for uml_moo, 71
 route command and routing, 57
 for bridging, 138
 for host setup, 91
 for instances, 76–77
 for multicast networks, 158
 for network interfaces, 255
 for TUN/TAP. *See* TUN/TAP devices
 for uml_switch, 165
 Running commands within instances, 180–182

S

- s command in sysrq, 174
- s switch
 - for screen, 204
 - for ubd, 303
 - for uname, 168
- Scaling in application-level clustering, 294
- Schedulers, guest, 295–297
- schedulers file, 295–297
- screen tool, 204–205
- Searching file contents, 114
- Secure mod_perl, 283–285
- Security
 - in application-level clustering, 293
 - in bridging, 140
 - in captive UML, 288–289
 - for host setup, 93
 - for large servers, 211
 - configuration for, 212–216
 - console, 223–225
 - jailing instances, 216–223
 - skas3 vs. skas0, 225–226
 - for TUN/TAP devices, 129–132
- Seekable host files, 53
- Sending interrupts to instances, 179
- Separate kernel address space mode. *See* skas (separate kernel address space) mode
- Serial Line IP (SLIP)
 - configuring, 150
 - for frames, 54
 - for host network access, 144
 - for host setup, 88
- Serial lines
 - setting up, 40–47
 - virtual, 79–82
- Server consolidation, 8–10
- servers. *See* Large server management; Small server setup
- setuid files, 113–114
- shadow_fs metadata format, 111–113
- shadowfs file, 112
- Shared memory for device emulation, 82
- showconfig option, 305
- Shutting down, 59–60
- SIGBUS signal, 207
- SIGHUP signal, 189
- SIGINT signal, 179, 189, 203
- SIGIO signal, 36
- SIGKILL signal, 173
- Signals for instance management, 188–189
- SIGSEGV signal, 195, 200–201
- SIGTERM signal, 173, 189
- SIGTSTP signal, 203
- SIGWINCH signal, 36
- Simulating hardware, 273–274
- Single-System Image (SSI) cluster, 292–293
- Size
 - of backing files, 69
 - in copying data into instances, 86
 - of COW files, 65–66
 - of filesystems, 118
- skas (separate kernel address space) mode, 194–196
 - enabling, 240
 - skas0, 195–196
 - with jails, 218
 - vs. skas3, 225–226
 - working with, 200–201
 - skas3, 195–196
 - with jails, 218
 - vs. skas0, 225–226
 - working with, 198–200
 - for threads, 179
- Slave UMLs, 80–81
- SLIP (Serial Line IP)
 - configuring, 150
 - for frames, 54
 - for host network access, 144
 - for host setup, 88
- Slirp networking emulator
 - configuring, 150–151
 - for host network access, 144–145
- Small server setup, 191–192
 - execution modes in, 194–196
 - patches for, 201–202
 - skas0, 200–201
 - skas3, 198–200
 - tt, 197–198
 - Vanderpool and Pacifica, 202–203
 - kernel version in, 192–194
 - long-lived instances in, 203–205
 - for memory, 206–208
 - for networking, 206
 - recommendations for, 209–210
 - umid directories in, 209

- SMP (Symmetric Multi-Processing), 197
 - SMP option, 243–244
 - Sniffers, 146
 - Sockets
 - with hostfs, 107–108
 - with humfs, 110
 - in MConsole, 185
 - notifications with, 188
 - with uml_switch, 152–153
 - Solaris zones, 3
 - Source, downloading, 234–235
 - sparse switch, 175
 - Specialized configurations, 251
 - clusters, 265–273
 - large numbers of devices
 - memory, 257–265
 - network interfaces, 251–257
 - Spoofing in bridging, 142
 - ssh command, 94–95
 - ssh keys, 181
 - ssh sessions, random numbers for, 246
 - SSI (Single-System Image) cluster, 292–293
 - SSL option, 248
 - SSL_CHAN option, 248
 - Stack, dumping, 174, 178
 - stack command, 177–179, 310
 - Standard application programming interface
 - in captive UML, 289
 - STATIC_LINK option, 240–241
 - Static linking
 - configuration option for, 240–241
 - with jails, 218
 - stop command, 310
 - Stopping
 - instances, 174–175
 - virtual machines, 117
 - strace tool, 226
 - su with jails, 220
 - Subnets for instances, 76
 - Subsystems, virtualized, 295–298
 - superblock files, 111–112, 118
 - Swap space
 - adding, 47–49
 - with compartments, 297
 - for instances, 26
 - for jails, 223
 - MADV_TRUNCATE for, 228
 - partitions for, 52
 - performance of, 259–265
 - swapoff command, 223
 - swapon command, 52, 223
 - switch command, 309
 - switch-tap option, 153
 - Switches
 - for packets, 72
 - virtual, 136
 - Symbolic links, 113
 - Symmetric Multi-Processing (SMP), 197
 - Synchronization
 - clock, 28
 - in sysrq, 174
 - Synchronous files, 115–117
 - systemu patch, 226–227
 - sysrq command and SysRq handler
 - for backups, 116–117
 - invoking, 172–174
 - MAGIC_SYSRQ for, 246
 - in uml_mconsole, 310
 - sysrq file, 172
 - System call tracing, 212
 - System-level programming, 12
 - System memory savings, COW files for, 64
- ## T
- t command in sysrq, 174
 - t switch
 - for iptables, 92
 - for tuncctl, 124, 311
 - for xterm, 302
 - tap switch for uml_switch, 311
 - Tape drives for copying data into instances, 84–85
 - tar files
 - for copying data into instances, 84–85
 - copying into UML, 53
 - length of, 87
 - for moving backing files, 69
 - for source, 235
 - Tasks
 - killing, 173
 - real-time, 174
 - tcpdump
 - for multicast networks, 158–159
 - with pcap, 145
 - for TUN/TAP, 125–127, 131–132
 - for uml_switch, 164–165

- telnet, 43, 45–47
 - telnetd, 46
 - Terminal emulators, 302
 - Testing
 - COW files for, 64
 - testbeds for, 9–10
 - TUN/TAP devices, 135–136
 - Threads
 - in execution modes. *See* Execution modes
 - forcing into contexts, 177–179
 - 3_LEVEL_PGTABLES option, 245
 - Time to live (TTL) setting, 151–152
 - Timers
 - listing, 35–36
 - real-time clock for, 245–246
 - Timestamps, 69–70
 - Timing bugs, 13
 - tmp directory
 - for bind mounts, 214–215
 - copying files to, 67, 82
 - for databases, 281
 - for filesystems, 34, 103
 - for jails, 219, 221
 - for memory, 31–32, 206–207
 - for processes, 106
 - tmpfs filesystem, 32, 34, 102, 207
 - Torvalds, Linus, 7
 - touch command, 70
 - Tracing thread (tt) mode, 194–196
 - enabling, 240
 - for threads, 178
 - working in, 197–198
 - Traffic analysis tools, 146
 - Translation of filesystem requests, 104
 - Transports
 - configuration options for, 247–248
 - networking, 142–143
 - configuring, 147–154
 - for host network access, 143–145
 - for isolated networks, 145–146
 - selecting, 146–147
 - tt (tracing thread) mode, 194–196
 - enabling, 240
 - for threads, 178
 - working in, 197–198
 - TTL (time to live) setting, 151–152
 - TTY_CHAN option, 248
 - tty_log_dir option, 305
 - tty_log_fd option, 305
 - tun file, 55, 57, 123
 - TUN/TAP devices, 35, 57
 - bridging with, 136–142
 - for frames, 54
 - for host network access, 143
 - in host setup, 88–90
 - with routing, 121–122
 - configuring, 122–124, 149
 - connectivity in, 125–129
 - DHCP for, 134–135
 - for outside network access, 132–133
 - security for, 129–132
 - testing, 135–136
 - tunctl utility
 - reference, 310–311
 - working with, 122–124
 - tuntap command, 149
- ## U
- u command in sysrq, 174
 - u switch for tunctl, 124, 311
 - ubd devices, 34
 - advantages of, 119
 - for filesystem access, 114–115
 - image backup for, 117
 - partitioning, 50
 - ubd option, 303
 - ubd0 file, 26
 - ubda switch for COW files, 62
 - ubdb file, 26
 - for copying data into instances, 84–85
 - for swap space, 48
 - ubd option, 303
 - UID root, 109
 - UIDs (user IDs)
 - in filesystem extensions, 118
 - in ownership, 106–107
 - umid (unique machine ID), 42
 - umid directory
 - with jails, 219
 - process IDs in, 189
 - in small server setup, 209
 - umid option, 304
 - uml.ctl socket, 152
 - uml_dir option, 305

uml_mconsole command, 182–183. *See also*
 Management Console (MConsole)
 for bridging, 138
 for copying data into instances, 84
 for devices, 43
 for host ports, 41–45
 for hosts, 87, 89
 for network devices, 56
 reference, 308–310
 for TUN/TAP device connectivity, 125
 for virtual serial lines, 80
 uml_moo tool
 for merging COW files, 71
 reference, 308
 UML_NET options, 247–248
 uml_net utility, 125, 129–130, 312
 UML_RANDOM option, 246–247
 UML_REAL_TIME_CLOCK option,
 245–246
 uml_switch process, 148
 configuring, 152–154
 example, 160–166
 for isolated networks, 145
 reference, 311
 UML_WATCHDOG option, 247, 312
 uname command, 168
 Unique machine ID (umid), 42
 Unique machine id (umid) directory
 with jails, 219
 process IDs in, 189
 in small server setup, 209
 UNIX sockets
 with hostfs, 107–108
 with humfs, 110
 in MConsole, 185
 with uml_switch, 152–153
 -unix switch for uml_switch, 311
 Unplugging devices, 169–170
 Unprivileged contexts, 231
 untar command, 53, 85
 User IDs (UIDs)
 in filesystem extensions, 118
 in ownership, 106–107
 User mode vs. kernel mode, 212
 User notifications, 188
 Utilities reference
 humfsify, 307–308

internal, 312
 tunctl, 310–311
 uml_mconsole, 308–310
 uml_moo, 308
 uml_switch, 311

V

-v option for uname, 168
 valgrind, 244
 Vanderpool execution mode, 202–203
 VCPU, 231
 version command, 168–169, 309
 --version option, 305
 Version queries, 168–169
 Virtual filesystems, 101–104
 Virtual hardware configuration options,
 246–247
 Virtual machines, purpose of, 3–4
 Virtual memory, 230–231
 Virtual operating systems, 2
 Virtual Private Networks (VPNs), 89
 Virtual processes, 282–283
 Virtual processors, 243–244
 Virtual serial lines, 79–82
 Virtual switches, 136
 Virtualized subsystems, 295–298
 vmalloc space, 258
 vmlinux file, 249–250
 vmlinuz file, 249
 VMWare technology, 2–3
 VPNs (Virtual Private Networks), 89
 vserver project, 2
 vtund, 89

W

WATCHDOG option, 247
 Web site requests, 284–285
 wget
 for connectivity, 74
 for host setup, 94
 winch interrupt, 36
 Wright, Chris, 234
 Write-protecting hostfs directories, 216

X

- X11 utilities, 96
- Xauthority application, 95
- .Xauthority file, 95
- xconfig configurator, 235–237
- xdpyinfo, 96
- Xen technology, 2–3
- xhost application, 95, 98
- xload, 96

- Xnest, 96–99
- xorg-x11-tools package, 96n
- xterm option, 43–44, 96, 302
- XTERM_CHAN option, 248
- xterm windows, 29

Z

- Zones, Solaris, 3